arXiv:1611.01929v2 [cs.AI] 8 Nov 2016

Deep Reinforcement Learning with Averaged Target DQN

Oron Anschel, Nir Baram, Nahum Shimkin
Department of Electrical Engineering
Technion - Israel Institute of Technology
Haifa 23000, Israel
{oronanschel@campus, nirb@campus, shimkin@ee}.technion.ac.il

Abstract

The commonly used Q-learning algorithm combined with function approximation induces systematic overestimations of state-action values. These systematic errors might cause instability, poor performance and sometimes divergence of learning. In this work, we present the AVERAGED TARGET DQN (ADQN) algorithm, an adaptation to the DQN class of algorithms which uses a weighted average over past learned networks to reduce generalization noise variance. As a consequence, this leads to reduced overestimations, more stable learning process and improved performance. Additionally, we analyze ADQN variance reduction along trajectories and demonstrate the performance of ADQN on a toy Gridworld problem, as well as on several of the Atari 2600 games from the Arcade Learning Environment.

1 Introduction

In Reinforcement Learning (RL) an agent seeks to find an optimal policy for a sequential decision problem (Sutton and Barto, 1998). One of the most popular algorithm for RL problems is the Q-learning algorithm (Watkins and Dayan, 1992). Here the agent learns sequential estimates of the optimal state-action value function:

$$Q_*(s,a) \triangleq \max_{\pi} \mathbb{E}^{\pi}[\sum_{t=0}^{\infty} \gamma^t R_t | s_0 = s, a_0 = a]$$

For small problems the Q-learning algorithm can be implemented using a table containing current value estimates of all possible state-action pairs, and will benefit from convergence guarantees under mild assumptions. In practice, most of the real-world problems are far too complex for a tabular implementation and require the use of function approximation.

One of the most impressive early works, bringing together RL and neural networks presented a world-class backgammon agent (Tesauro, 1992). Nevertheless, combined directly, divergence was shown even on simple problems (Boyan and Moore, 1995), and until recent breakthroughs the progress in the field was slow.

The DQN algorithm (Mnih et al. (2013)) combines Q-learning with the recent progress in neural network techniques. DQN was the first successful attempt to learn control policies directly from a high-dimensional sensory input, and led the way for many subsequent developments in the field. Two significant modifications were introduced in the DQN work that help to stabilize the learning process: First, the network is trained using samples from an experience-buffer. Second, the target values are calculated using a target network whose parameters are frozen for a fixed number of iterations at a time.

Although DQN proved to be successful in most of the benchmark problems (Mnih et al., 2013), in some of them it seems to fail. One of the reasons for the failure of DQN is argued to be insufficient exploration strategy, such as in the Atari game MONTEZUMA'S REVENGE. Several works suggested

exploration techniques for high-dimensional state space to remedy this problem (Bellemare et al., 2016; Osband et al., 2016).

Another possible reason for the DQN to fail in Atari games such as WIZARD OF WOR and ASTERIX is systematic errors caused by the use of Q-learning combined with function approximation. Essentially, high variance in target values combined with the max operator might lead to a positive bias, which at the extreme might leads to divergence. In this work, we propose the ADQN algorithm as a solution to this problem. Additionally, we analyze and show how a straightforward approach which we denote as *Ensemble* DQN is inferior in variance reduction.

The rest of the paper is constructed as follows. In Section 2 we discuss relevant related work. Section 3 elaborates on the Q-learning and DQN algorithms. In section 4 we present the ADQN algorithm. In Section 5 we discuss and analyze errors bias and variance induced by the combination of function approximation with Q-learning. Section 6 provides an empirical evaluation of the ADQN algorithm both in a toy problem and in several of the Atari 2600 games from the Arcade Learning Environment (Bellemare et al., 2013). Section 7 closes the paper with our conclusions and possible future direction.

2 Related Work

Thrun and Schwartz (1993) is an early theoretical work analyzing the effects of combining function approximation and RL techniques. This work modeled the generalization noise induced by the use of approximator as zero-mean independent random variables and showed that systematic overestimation errors could occur due to the utilization of the max operator in the standard Q-learning update rule.

Double Q-learning (Hasselt, 2010) was suggested as a solution for this problem. The max operator was divided into selection and evaluation parts, which replaced the positive bias of the original predictor with a negative one. Later the Double-DQN (DDQN) algorithm that was introduced in Van Hasselt et al. (2015) provided a DQN implementation of this algorithm.

Since the publication of the DQN work, many variants and modification of the standard algorithm have been developed. Schaul et al. (2015) showed improved results by employing smart policies for sampling from the experience buffer; Wang et al. (2015) introduced a modified Q-network architecture; Mnih et al. (2016) implemented asynchronous version of the DQN which does not require the use of an experience-buffer and showed impressive performance.

Another recent work, the Bootstrapped DQN algorithm (Osband et al., 2016) implemented *Deep Exploration* with the use of several Q-networks learning in parallel. Bootstrapped DQN is similar to our work by the utilization of an ensemble of several Q-networks. However, the primary purpose of Bootstrapped DQN is *exploration* whereas our objective is *variance reduction*. We note two additional differences between the Bootstrapped DQN and this work: (1) In Bootstrapped DQN the different Q-networks are used for gathering shared experience during the learning process, while in our algorithm they are combined continuously in order to construct target values; (2) In Bootstrapped DQN the networks learn in parallel while in our work there is a single learning network.

We note that the variants and modification mentioned above are essentially orthogonal to our work, so that combining them with our proposed scheme is likely to yield further performance improvements.

3 Deep Q Networks (DQN)

In classical Q-learning (Watkins and Dayan (1992)) the update rule is given by

$$Q_{t+1}(s,a) \leftarrow (1-\alpha)Q_t(s,a) + \alpha(r(s,a) + \gamma \max_{a'} Q_t(s',a')) \tag{1}$$

where s' is the resulting state after applying action a in the current state s, γ is the MDP's discount factor and α is a time varying learning rate. Under mild assumption for finite state and action spaces, it can be shown that the Q-learning algorithm converges to the optimal action-value function, as long as every state action pair is updated infinitely often. The update of the Q-function estimate can be done in an online or offline manner.

The above Q-learning update rule can be directly implemented in a neural network. Since no direct assignment of Q-values can be performed as if we had a table of state action pairs, a loss function

must be introduced (Riedmiller, 2005). In case the squared error is taken as a loss function (Eq. 3), gradient descent takes the form of Eq. 4:

$$y_t^Q = \mathbb{E}_{s' \sim \mathcal{B}} \left[r(s,a) + \gamma \max_{a'} Q(s', a'; \theta^-) | s, a \right] \tag{2}$$

$$L(\theta_t) = \mathbb{E}_{(s,a) \sim \mathcal{B}} \left[(y_t^Q - Q(s,a;\theta_t))^2 \right] \tag{3}$$

$$\theta_{t+1} \leftarrow \theta_t + \eta \mathbb{E}_{(s,a) \sim \mathcal{B}} \left[(y_t^Q - Q(s,a;\theta_t)) \nabla_\theta Q(s,a;\theta_t) \right] \tag{4}$$

Here η is a scalar learning rate, y_t^Q is the target value with $\theta^- = \theta_t$, and $\mathcal{B}$ is an experience buffer which normally contains one tuple of experience $\mathcal{B} = \{s_t, a_t, s_{t+1}, r_t\}$. The DQN algorithm (Mnih et al., 2013) introduced two important modifications for the algorithm above. First, the algorithm samples from a large experience buffer $\mathcal{B}$ which contains many experience tuples. Second, the algorithm uses a target network with parameters θ^- that are copied from the learning network $\theta^- \leftarrow \theta_t$ only every τ iterations. These last modifications enable the Q-learning algorithm to operate in a semi off-policy manner, breaking the task of learning the optimal state-action values into smaller supervised learning problems.

4 Averaged Target DQN (ADQN)

We introduce here the ADQN algorithm. ADQN maintains a target networks buffer which contains the parameters of K past learned networks. Also, an additional learning rate parameter $\alpha \in [0, 1)$ is required. Essentially the DQN algorithm can be viewed as a special case of the ADQN algorithm with a target networks buffer of size one. The only algorithmic difference from the DQN algorithm is the different choice of target value as in Eq. 5. The loss function (Eq. 6) remains the same as in DQN and so does the parameter update equation (Eq. 7):

$$y_t^A = \mathbb{E}_{s' \sim \mathcal{B}} \left[r(s,a) + \gamma \max_{a'} \sum_{k=0}^{K-1} w_k Q(s', a'; \theta_{\tau(\lfloor t/\tau \rfloor - k)}) | s, a \right] \tag{5}$$

$$L(\theta_t) = \mathbb{E}_{(s,a) \sim \mathcal{B}} \left[(y_t^A - Q(s,a;\theta_t))^2 \right] \tag{6}$$

$$\theta_{t+1} \leftarrow \theta_t + \eta \mathbb{E}_{(s,a) \sim \mathcal{B}} \left[(y_t^A - Q(s,a;\theta_t)) \nabla_\theta Q(s,a;\theta_t) \right] \tag{7}$$

where $w_k = (1-\alpha)^k / C_w$ and $C_w = \sum_{k=0}^{K-1} w_k$ is a normalization constant. Every τ iterations, the parameters of the learning networks are copied to the target networks buffer and the oldest parameters are removed. In Eq. 5 the summation is performed over the K networks in the target buffer. The learned Q-function at time t is the averaged network:

$$Q^A(s,a;\bar{\theta}_t) = \sum_{k=0}^{K-1} w_k Q(s,a;\theta_{\tau(\lfloor t/\tau \rfloor - k)}) \tag{8}$$

The learning rate parameter α trades between variance reduction and valuing newer Q estimates. In general, α can be a function of time and even the state space.

5 Estimation Errors in RL

Combining function approximation and RL methods introduces difficulties that do not exist for each separately. Thrun and Schwartz (1993) provided a first theoretical analysis for combing Q-learning and function approximations. In this section, we will discuss their results and suggest further analysis of variance propagation along trajectories.

	$\mathbb{E}_z\left[(\bar{Q}(s_2,a)^2\right]$	$K=10$	$K=20$
DQN	$\sigma_z^2(1+\gamma^2)$	$\sigma_z^2(1+\gamma^2)$	$\sigma_z^2(1+\gamma^2)$
ADQN	$\sigma_z^2\left[\frac{1}{K}+\frac{\gamma^2}{K^4}\left(2\sum_{k=1}^{K-1}k^2+K^2\right)\right]$	$\sigma_z^2\left(\frac{1}{10}+\frac{\gamma^2}{14.92}\right)$	$\sigma_z^2\left(\frac{1}{20}+\frac{\gamma^2}{29.96}\right)$
Ensemble DQN	$\sigma_z^2\left(\frac{1+\gamma^2}{K}\right)$	$\sigma_z^2\left(\frac{1}{10}+\frac{\gamma^2}{10}\right)$	$\sigma_z^2\left(\frac{1}{20}+\frac{\gamma^2}{20}\right)$

Table 1: Learned Q-function variance in state s_2 of the MDP in Figure 1.

5.1 Overestimation

Thrun and Schwartz (1993) considered an additive zero-mean random variable, a generalization error model for the Q-function. That is, $Q(s',a';\theta^-)=Q^{target}(s',a')+Z_{s',a'}$ in Eq. 2. Due to the max operator in the equation, up to $\gamma\epsilon\frac{n-1}{n+1}$ overestimation in each target action-value $y_t^Q(s,a)$ can be seen on average (where $Z_{s,a}$ is uniformly distributed in an interval $[-\epsilon,\epsilon]$, and n is the number of applicable actions in state s). Systematic errors occur due to the nature of the Q-learning update equations, where along trajectories of the current greedy policy, each state can be overestimated up to $\gamma\epsilon\frac{n-1}{n+1}$ plus the bias of the next state in the trajectory. Hasselt (2010) argued that even in a tabular setting environment noise can cause overestimations and suggested the Double Q-learning algorithm as a solution.

We observe two adverse effects of large variance in the estimated Q-function: First, evidently, the learned Q-function is likely to deviate from the required one; Second, large variance in Q causes positive bias in the target values y_t^Q as well as variance which might lead to divergence.

5.2 Variance Reduction due to Averaging

In the following illustrative example, we analyze the noise variance propagation along trajectories, from the use of ADQN, DQN and an Ensemble technique that will be introduced later in this section.

Let $0,1,\ldots,t$ denote the indices of time instants in which parameters are copied from the learned network to the target buffer (DQN maintains a buffer of size one). We model the noise in the learned Q-function, in Eq. 2 or 5 as $Q(s',a';\theta_t)=y_{t-1}(s',a')+Z_{s',a'}$, where y_{t-1} is the previous time step target value. The noise $Z_{s,a}$ is modeled as a zero-mean random variable with variance σ_z^2, independent of the state. Also, we assume that τ, the number of iterations between time t and $t+1$ is large enough so that $Z_{s,a}^t$ and $Z_{s,a}^{t+1}$ are not correlated. The Q-networks themselves are not stochastic, and after the network parameters are copied to the target buffer the noise is also fixed. We are interested in calculating the variance of the learned policy which is defined as $\bar{Q}(s,a)\triangleq Q(s,a;\theta_t)$ for DQN, and $\bar{Q}(s,a)\triangleq Q^A(s,a;\bar{\theta}_t)$ for ADQN (Eq. 8).

Consider a two states, single action MDP such as in Figure 1. The process starts at state s_2, then moves to a terminal state s_1. A zero reward is obtained in both states.

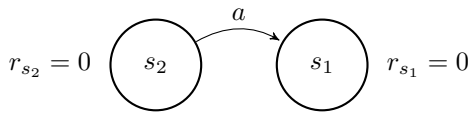

Figure 1: Two states, single action MDP. The process starts at state s_2, then moves to a terminal state s_1. A zero reward is obtained in both states.

Obviously, for all pairs (s,a), the optimal action-value function $Q^*(s,a)=0$, and therefore $\mathbb{E}_z\left[(\bar{Q}(s,a)-Q^*(s,a))^2\right]=\mathbb{E}_z\left[(\bar{Q}(s,a))^2\right]$. The analysis here reflects variance analysis for ADQN or DQN, assuming that the estimated Q-values are different enough along a current greedy policy trajectory so that the max operator in Eq. 2 can be ignored. The results are summarized in Table 1 (calculations are detailed in Appendix A).

ADQN takes an average over past learned Q-functions. One can consider a more direct approach to reducing variance. To reduce noise variance in the learned Q-networks we can learn several

	$\mathbb{E}_z\left[(\bar{Q}(s_L,a)^2\right]$
DQN	$\sigma_z^2 \sum_{l=1}^{L} \gamma^{2(l-1)}$
ADQN	$\sigma_z^2 \sum_{l=1}^{L} \gamma^{2(l-1)} \frac{D^{K,l}}{K^{2l}}$
Ensemble DQN	$\sigma_z^2 \sum_{l=1}^{L} \gamma^{2(l-1)} \frac{1}{K}$

Table 2: Learned Q-function variance in state s_L of the MDP in Figure 2.

independent estimates from the same target, and then average out the noise. This leads to an architecture of K independent Q-learners with distinct simulators and experience buffers, sharing only the learned policy while constructing target values. Denote the k-th network parameters at time t as θ_t^k. The learned Q-function in this architecture is the average over the K target networks estimates:

$$Q^E(s_2,a;\bar{\theta}_t) \triangleq \frac{1}{K}\sum_{k=0}^{K-1} Q(s',a';\theta_t^k)$$

and the target values for the k-th leaner at time t are computed as:

$$y_t^{E,k}(s,a) \triangleq \mathbb{E}_{s'\sim\mathcal{B}_k}\left[r(s,a)+\gamma \max_{a'}\frac{1}{K}\sum_{k=0}^{K-1} Q(s',a';\theta_{t-1}^k)|s,a\right]$$

Surprisingly, the above described architecture which we refer to as *Ensemble* DQN is not only computationally demanding (learning K networks in parallel instead of one), but also leads to inferior variance reduction as summarized in Table 1 (calculations are detailed in Appendix B).

The above calculations can be generalized to an L state, single action MDP as in Figure 2 (see Appendix B). ADQN's performance in variance reduction improves as the trajectory gets longer.

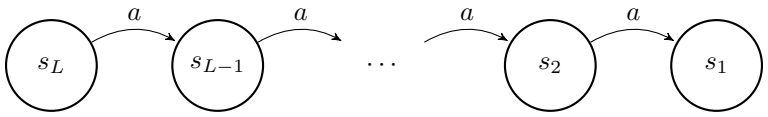

Figure 2: L states, single action MDP. The process starts at state s_L, then in each time step moves to the right, until the terminal state, state s_1 is reached. A zero reward is obtained in any state.

This is due to the time averaging which leads to average over averages and can be observed in the calculations. We denote by u^K a discrete rectangle function ($u_i^K = \mathbb{1}\{i \in (0,1,\ldots,K-1)\}$), and by $U^{K,DFT}$ its Discrete Fourier Transform (DFT). Also, we define $D^{K,L} \triangleq \frac{1}{N}\sum_{n=0}^{N-1}|U_n^{K,DFT}|^{2L}$ (where N is the length of $U^{K,DFT}$). Results, and the behaviour of the ADQN variance reduction coefficient $D^{K,L}/K^{2L}$, are summarized in Table 2 and Figure 3.

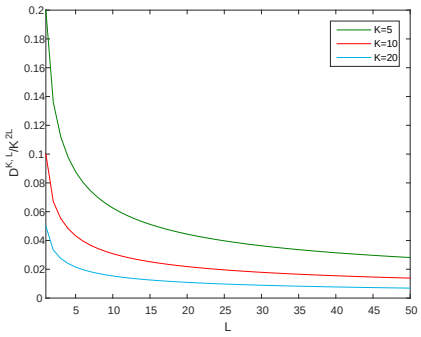

Figure 3: ADQN variance reduction coefficient $D^{K,L}/K^{2L}$ function. Note that for $L = 1$: $D^{K,L}/K^{2L} = 1/K$ which is the Ensemble DQN coefficient.

6 Experiments

We use two different environments to evaluate the performance of the ADQN algorithm and some assumptions made in previous sections. First, we will show experiments on a toy problem of Gridworld where the optimal value function can be computed exactly with a tabular state space representation. Second, we will demonstrate and compare performance using the Arcade Learning Environment (Bellemare et al. (2013)), which is the most commonly used benchmark environment for RL algorithms.

6.1 Gridworld

In this experiment on the problem of 2D Gridworld, the state space contains pairs of points from a 2D discrete grid ($S = \{(x, y)\}_{x,y\in 0,\dots,N}$). The algorithm interacts with the MDP through raw pixel features with a 2D one-hot feature map $\phi(s_t) := (\mathbb{1}\{(x, y) = s_t\})_{x,y\in 0,\dots,N}$. There are four actions corresponding steps in each compass direction, a reward of $r = 1$ if $s_t = (N, N)$ (else the reward is zero) and a discount factor of $\gamma = 0.9$.

Figure 4 shows the mean value estimation over all states for the ADQN algorithm with different size of target networks buffer. Also, we have plotted a baseline of the true optimal mean value V^*. On the left plot, we can see a negative effect of the slowdown in convergence in ADQN, and this is since newer estimation of the state-action values are more likely to be better than older ones. On the right, we have zoomed in to see how systematic overestimations cause positive bias, and how adding target networks reduces both the variance between the models and the overall bias. For the experiments, we have used the ADAM optimizer (Kingma and Ba (2014)), updated the target network parameters each $\tau = 300$ iterations, and $\alpha = 0.01$ was set as the ADQN learning rate.

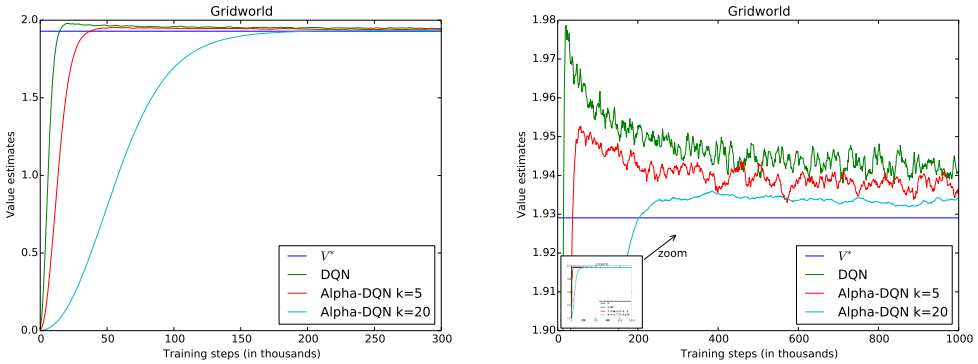

Figure 4: Comparison of value estimates during learning, comparing different values for the size of target networks buffer K, where for $K = 1$ the ADQN reduces to DQN. Results were averaged over 12 experiments.

In Figure 5 we have used the same setup and hyper-parameters as mentioned above. We have compared ADQN with $K = 20$ target networks to the Ensemble DQN method mentioned in Section 5.2. The Ensemble DQN learns in parallel 20 Q-networks for $\tau = 300$ iterations and then constructs the next target and test Q-network by averaging over the 20 networks that were learned. On the left plot, we can see that Ensemble DQN converges as fast as DQN, and ADQN suffers from a slowdown in convergence due to the use of a fixed small learning rate α. On the right plot we see that although Ensemble DQN is much more demanding computationally, ADQN converged to a lower bias faster and with less variance between successive models.

6.2 Arcade Learning Environment

We have evaluated the ADQN algorithm on several Atari games on the Arcade Learning Environment (Bellemare et al. (2013)). The Arcade Learning Environment is the most commonly used benchmark to evaluate DQN variants and adaptations, and this is due to the diverse set of games that need to be solved with a single set of architecture and hyper-parameters, directly from raw pixels input and

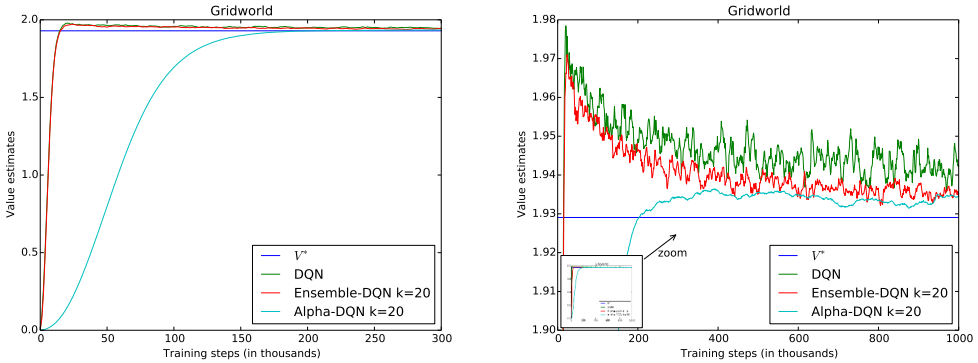

Figure 5: Comparison of value estimates during learning, comparing ADQN with $K = 20$ target network to ensemble DQN with 20 Q-networks learn in parallel averaged each $\tau = 300$ iterations to construct ensemble target. The DQN algorithm is presented as a reference point. Results were averaged over eight experiments.

a reward signal. We followed the setup of hyper-parameters and network architecture as in (Mnih et al. (2013)). Additionally, we have used moderate computation and memory architecture for ADQN with $K = 10$ target networks buffer and $\alpha = 0.01$. This configuration of the ADQN presents a 35% slowdown from the standard DQN, and the slowdown is due to the need of K times more forward computation (compared to DQN) to construct the next target values.

Figure 6 presents a visual comparison. ADQN shows improvement in the learned policies over DQN, the learning progress is more stable, and the value estimation is considerably more conservative. Also, the *spikes* phenomena (e.g., Wizard of Wor value estimate plot) that occurs in value estimates of DQN (and in Double-DQN) does not take place in the ADQN algorithm (which could be an indication of remedy to the systematic overestimation that occurs in the DQN algorithm).

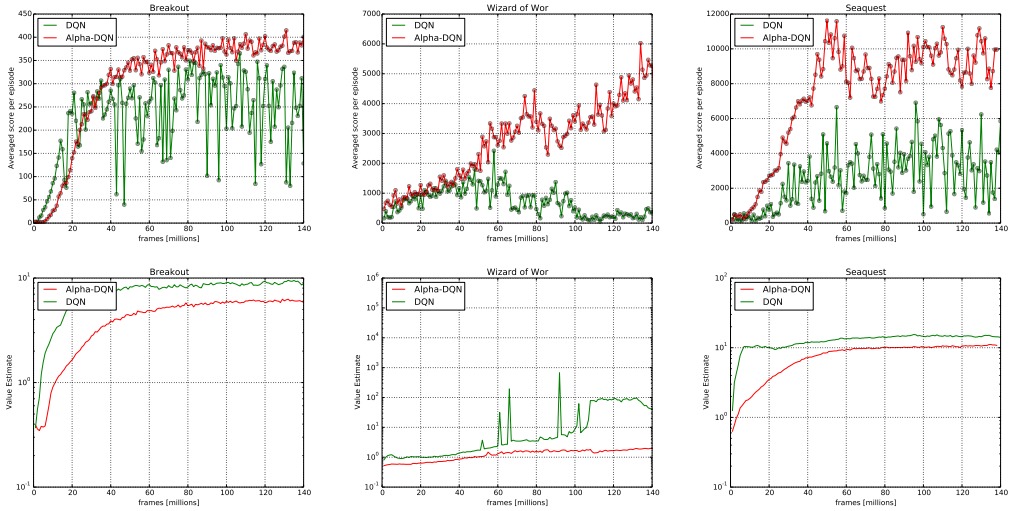

Figure 6: The top row shows the average reward per episode achieved by the algorithm during training evaluated with ϵ-greedy policy with $\epsilon = 0.05$ for 125,000 steps. The bottom row shows the average value estimate on a held-out set of states taken from the experience buffer. The stability of the ADQN, as well as the improved result and conservative value estimates can be seen.

7 Conclusions and Future Directions

In this work, we have presented the ADQN algorithm, an efficient variance reduction technique for the DQN family of algorithms. We have analyzed and empirically showed why a time-average approach is more beneficial for variance reduction, compared to a direct averaging approach such as the Ensemble DQN mentioned above. Using the Arcade Learning Environment we have shown how the learning process is stabilized along with improvement in performance, and the *spikes* phenomena that occur in both DQN and DDQN for some games, does not take place anymore.

In future work, it may be interesting to see how α, the ADQN learning rate parameter can be learned as a function of the states or features, and to what extent combining the proposed averaging scheme with other techniques will prove beneficial. Furthermore, as our variance analysis predicts, the ADQN technique is likely to be beneficial in value estimation schemes, and further investigation of the subject is required. Finally, the average is just one of the ensemble techniques that could be considered. Other include a majority vote, median, etc., over the learned Q-functions, and in the target values computations.

References

Bellemare, M. G., Naddaf, Y., Veness, J., and Bowling, M. (2013). The arcade learning environment: An evaluation platform for general agents. *Journal of Artificial Intelligence Research*, 47:253–279.

Bellemare, M. G., Srinivasan, S., Ostrovski, G., Schaul, T., Saxton, D., and Munos, R. (2016). Unifying count-based exploration and intrinsic motivation. *arXiv preprint arXiv:1606.01868*.

Boyan, J. and Moore, A. W. (1995). Generalization in reinforcement learning: Safely approximating the value function. *Advances in neural information processing systems*, pages 369–376.

Hasselt, H. V. (2010). Double Q-learning. In Lafferty, J. D., Williams, C. K. I., Shawe-Taylor, J., Zemel, R. S., and Culotta, A., editors, *Advances in Neural Information Processing Systems 23*, pages 2613–2621. Curran Associates, Inc.

Kingma, D. P. and Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv: 1412.6980*.

Mnih, V., Badia, A. P., Mirza, M., Graves, A., Lillicrap, T. P., Harley, T., Silver, D., and Kavukcuoglu, K. (2016). Asynchronous methods for deep reinforcement learning. *arXiv preprint arXiv:1602.01783*.

Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., and Riedmiller, M. (2013). Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*.

Osband, I., Blundell, C., Pritzel, A., and Van Roy, B. (2016). Deep exploration via bootstrapped dqn. *arXiv preprint arXiv:1602.04621*.

Riedmiller, M. (2005). Neural fitted q iteration–first experiences with a data efficient neural reinforcement learning method. In *European Conference on Machine Learning*, pages 317–328. Springer.

Schaul, T., Quan, J., Antonoglou, I., and Silver, D. (2015). Prioritized experience replay. *arXiv preprint arXiv:1511.05952*.

Sutton, R. S. and Barto, A. G. (1998). *Reinforcement Learning: An Introduction*, volume 1. MIT Press Cambridge.

Tesauro, G. (1992). Practical issues in temporal difference learning. In *Reinforcement Learning*, pages 33–53. Springer.

Thrun, S. and Schwartz, A. (1993). Issues in using function approximation for reinforcement learning. In *Proceedings of the 1993 Connectionist Models Summer School Hillsdale, NJ. Lawrence Erlbaum*.

Van Hasselt, H., Guez, A., and Silver, D. (2015). Deep reinforcement learning with double Q-learning. *arXiv preprint arXiv: 1509.06461*.

Wang, Z., de Freitas, N., and Lanctot, M. (2015). Dueling network architectures for deep reinforcement learning. *arXiv preprint arXiv: 1511.06581*.

Watkins, C. J. and Dayan, P. (1992). Q-learning. *Machine Learning*, 8(3-4):279–292.

A Variance Calculation for the Two-State Model in Section 5.2

We outline here the calculation of variances that were quoted in Section 5.2. We consider the DQN, ADQN and Ensemble DQN.

DQN:

$$\begin{aligned}
y^Q_{t-2}(s_1,a) &= 0 \\
Q(s_1,a;\theta_{t-1}) &= Z^{t-1}_{s1} + y^Q_{t-2}(s_1,a) \\
&= Z^{t-1}_{s1} \\
y^Q_{t-1}(s_2,a) &= 0 + \gamma Q(s_1,a;\theta_{t-1}) \\
Q(s_2,a;\theta_t) &= Z^t_{s_2} + y^Q_{t-1}(s_2,a) \\
&= Z^t_{s_2} + \gamma Q(s_1,a;\theta_{t-1}) \\
&= Z^t_{s_2} + \gamma Z^{t-1}_{s_1} \\
\mathbb{E}_z\left[(\bar{Q}(s_2,a))^2\right] &= \mathbb{E}_z\left[(Q(s_2,a;\theta_t))^2\right] \\
&= \sigma_z^2(1+\gamma^2)
\end{aligned}$$

ADQN with K target networks buffer:

$$\begin{aligned}
y^A_{t-2}(s_1,a) &= 0 \\
Q(s_1,a;\theta_{t-1}) &= Z^{t-1}_{s_1,a} + y^A_{t-2}(s_1,a) \\
&= Z^{t-1}_{s_1,a} \\
y^A_{t-1}(s_2,a) &= 0 + \gamma\left[\frac{1}{K}\sum_{k'=0}^{K-1} Q(s_1,a;\theta_{t-1-k'})\right] \\
&= 0 + \gamma\left[\frac{1}{K}\sum_{k'=0}^{K-1} Z^{t-1-k'}_{s_1,a}\right] \\
Q(s_2,a;\theta_t) &= Z^t_{s_2} + y^A_{t-1}(s_2,a) \\
Q^A(s_2,a;\bar{\theta}_t) &= \frac{1}{K}\sum_{k'=0}^{K-1} Q(s_2,a;\theta_{t-k'}) \\
&= \frac{1}{K}\sum_{k'=0}^{K-1}\left[Z^{t-k'}_{s_2} + y^A_{t-k'-1}(s_2,a)\right] \\
&= \frac{1}{K}\sum_{k'=0}^{K-1}\left[Z^{t-k'}_{s_2}\right] + \frac{1}{K}\sum_{k'=0}^{K-1}\left[\gamma\frac{1}{K}\sum_{m=0}^{K-1} Z^{t-k'-m-1}_{s_1}\right] \\
\mathbb{E}_z\left[Q^A(s_2,a;\bar{\theta}_t))^2\right] &= \sigma_z^2\left[\frac{1}{K} + \frac{\gamma^2}{K^4}\left(2\sum_{k=1}^{K-1} k^2 + K^2\right)\right]
\end{aligned}$$

Ensemble-DQN with K networks learn in parallel from shared leaned Q-function:

$$
\begin{aligned}
y_{t-2}^{E,k}(s_1,a) &= 0 \\
Q(s_1,a;\theta_t^k) &= Z_{s_1,a}^{t,k} \\
y_{t-1}^{E,k}(s_2,a) &= 0+\gamma\left[\frac{1}{K}\sum_{i=0}^{K-1} Q(s_1,a;\theta_{t-1}^k)\right] \\
&= 0+\gamma\left[\frac{1}{K}\sum_{i=0}^{K-1} Z_{s_1,a}^{t-1,k}\right] \\
Q(s_2,a;\theta_t^k) &= Z_{s_2,a}^{t,k}+y_{t-1}^{E,k}(s_2,a) \\
Q^E(s_2,a;\bar{\theta}_t) &= \frac{1}{K}\sum_{k=0}^{K-1} Q(s_2,a;\theta_t^k) \\
&= \frac{1}{K}\sum_{k=0}^{K-1} Z_{s_2,a}^{t,k}+y_{t-1}^{E,k}(s_2,a) \\
&= \frac{1}{K}\sum_{k=0}^{K-1}\left[Z_{s_2,a}^{t,k}+\gamma\left(\frac{1}{K}\sum_{j=0}^{K-1} Z_{s_1,a}^{j,t-1}\right)\right] \\
&= \frac{1}{K}\sum_{i=0}^{K-1}\left[Z_{s_2,a}^{t,k}+\gamma Z_{s_1,a}^{k,t-1}\right] \\
\mathbb{E}_z\left[(Q^E(s_2,a;\bar{\theta}_t))^2\right] &= \sigma_z^2\left(\frac{1+\gamma^2}{K}\right)
\end{aligned}
$$

B Trajectory Calculations for Section 5.2

We use the following notation for the Discrete Fourier Transform (DFT) in our calculations: For a vector x, $X_k^{DFT}=\sum_{j=0}^{K-1} x_j W_{K-1}^{jk}$ where $W_{K-1}=e^{-2\pi i/(K-1)}$ and $i=\sqrt{-1}$. We denote the discrete rectangle function as u^K which is $u_n^K=\mathbb{1}\{n\in(0,\ldots,K-1)\}$ and its DFT as $U^{DFT,K}$. Since Z^t are assumed to be zero-mean and independent random variables, it is easy to check that $\mathbb{E}\left[\left(\sum_{k_L=0,\ldots,k_2=0,k_1=0}^{K-1} Z^{k_1-k_2\ldots-k_L}\right)^2\right]=\sigma_z^2\sum|u_1*u_2\ldots*u_L|^2$ where $*$ is a circular convolutional and u_l are all discrete rectangle functions with support of K. Using Parseval's theorem we get that $\sum|u_1*u_2\ldots*u_L|^2=\sum_{n=0}^{N-1}|U_n^{DFT,K}|^{2L}/N$ where N is the actual length of the vector u that might need to be zero padded (Nyquist condition).